UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/691,515	10/24/2003	David James Wilson	3458-Z	7405

7590 06/21/2007

Law Office of Jim Zegeer
Suite 108
801 North Pitt Street
Alexandria, VA 22314

EXAMINER	
WEINTROP, ADAM S	
ART UNIT	**PAPER NUMBER**
2145	

MAIL DATE	DELIVERY MODE
06/21/2007	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No.	Applicant(s)	
	10/691,515	WILSON, DAVID JAMES	
	Examiner	Art Unit	
	Adam S. Weintrop	2145	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *24 October 2003*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-23* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-23* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☒ The drawing(s) filed on *24 October 2003* is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☒ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date *5/23/05*.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

DETAILED ACTION

Claim Objections

1. **Claims 1-23** are objected to because of the following informalities:

Regarding **claim 1**, the term "extension headers" in claim lines 5-6 has already been defined and should be replaced with --said extension headers-- to clarify the claim language. The term "upper-layer header" on claim line 7 is singular and should be replaced with --upper-layer headers-- to improve the clarity of the claim.

Regarding **claim 4**, the term "the first extension header" on claim lines 1-2 has not been defined and should be replaced with --a first extension header-- to improve the clarity of the claim language. The term "the group of extension headers" on claim line 2 should be replaced with --a group of extension headers-- to improve the clarity of the claim.

Regarding **claim 5**, the term "cache entry" on claim line 1 has already been defined and should be replaced with --said cache entry--.

Regarding **claim 7**, the term "upper-layer header" on claim line 2 is singular and should be replaced with --upper-layer headers-- to improve the clarity of the claim. The term "a cache lookup" on claim lines 2-3 has already been defined and should be replaced with --said cache lookup--.

Regarding **claim 14**, the term "upper-layer header" on claim lines 1-2 is singular and should be replaced with --upper-layer headers-- to improve the clarity of the claim.

Regarding **claim 17**, the term "extension headers" in claim line 5 has already been defined and should be replaced with --said extension headers-- to clarify the claim language. The term "upper-layer header" on claim line 6 is singular and should be replaced with --upper-layer headers-- to improve the clarity of the claim.

Regarding **claim 19**, the term "a cache lookup" on claim lines 1-2 has already been defined and should be replaced with --said cache lookup--.

Appropriate correction is required.

Claim Rejections - 35 USC § 103

2. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

3. **Claims 1-10 and 16-21** are rejected under 35 U.S.C. 103(a) as being unpatentable over Soirinsuo et al. (US 6,084,855) in view of "A proposal for the Ipv6 Flow Label" (Conta et al.).

Regarding **claims 1 and 17**, Soirinsuo et al. anticipates:

A method or a system of accessing upper-layer headers in a packet flow (column 7, lines 13-16, where L3 headers are accessed), comprising the

steps of:

a) responsive to a packet header containing extension headers, building a cache key and performing a cache lookup for a cache entry (column 7, line 64-column 8, line 14, where the packet containing extension headers is processed and the fields cached for subsequent packets); and

b) responsive to finding a corresponding cache entry, reading extension headers using the cache entry to arrive at and read fields in the header (column 7, line 64-column 8, line 14, where the cache entry is based on the processing of extension headers, seen as reading the extension headers using the cache entry to read fields in the header).

Soirinsuo et al. does not disclose reading the extension header in parallel or using the cache to retrieve upper layer header information. The general concept of reading extension headers in parallel and then reading the upper-layer header information is well known in the art as illustrated by Conta et al.

Conta et al. describes a packet classification system where the extension headers are read all at once (page 26, section A.3, where the lengths of the headers and extension headers are stored to enable the classifier to skip over the headers all at once, seen as a parallel read since they are not processed independently), and then the upper-layer header information is read and classified (page 26, section A.3, where after the header length information is used, the classifier can access ports, addresses, and protocol information).

It would have been obvious to one of ordinary skill in the art at the time of invention to modify Soirinsuo et al. with reading extension headers in parallel and retrieving upper-layer header information as taught by Conta et al. in order to speed up processing of Ipv6 packets by reducing the amount of information to be processed.

Regarding **claims 2 and 18**, Soirinsuo et al. and Conta et al. teach all of the limitations as described above, with Soirinsuo et al. further teaching:
The method as defined in claim 1 or the system in claim 17, wherein the packet flow comprises IPv6 packets (column 7, lines 64-67).

Regarding **claim 3**, Soirinsuo et al. and Conta et al. teach all of the limitations as described above, with Soirinsuo et al. teaching:
The method as defined in claim 2 wherein the cache lookup is performed on extension headers and the fields present in the IPv6 header (column 7, line 64-column 8, lines 14, where the processing is performed using the extension headers and the Ipv6 header, and the processing is cached, seen as using the cache lookup on extension headers and fields in the Ipv6 header).

Soirinsuo et al. does not teach using the cache lookup with a table containing lengths of the extension headers and then using a tuple based on the fields in the Ipv6 header.

The general concept of using a table of lengths of extension headers for classification and then using a tuple based on fields in the Ipv6 header is well known in the art as illustrated by Conta et al. Conta et al. teaches that a field of extension header lengths is used as part of the classification process (page 26, section A.3, where the classification process uses the field of extension header lengths, seen as a table, to classify the packet). Conta et al. also teaches that the classification can be performed using fields in the Ipv6 header such as port, address, and protocol information (page 26, section A.3, where the classifier can access the header information).

It would have been obvious to one of ordinary skill in the art at the time of invention to modify Soirinsuo et al. with reading extension header lengths in the classification system and retrieving upper-layer header information to be used in the classification system as taught by Conta et al. in order to speed up processing of Ipv6 packets by reducing the amount of information to be processed.

Regarding **claim 4**, Soirinsuo et al. and Conta et al. teach all of the limitations as described above, with Soirinsuo et al. teaching that classifications of packets can be processed and cached and retrieved with cache lookups (column 8, lines 1-14). Soirinsuo et al. does not teach reading the first extension header in the group of extension headers while the cache lookup is being performed. The general concept of reading a first extension header in a group of extension headers in the process of classification (which can be cached as Soirinsuo et al. teaches) is well known in the art

as illustrated by Conta et al. Conta et al. teaches that the header length header is read to determine the length of the rest of the extension header (page 26, section A.3, where the header length header field is seen as an extension header read first in the group as classification is being performed). It would have been obvious to one of ordinary skill in the art at the time of invention to modify Soirinsuo et al. with reading extension header lengths as the first extension header before the rest of the group in a classification system as taught by Conta et al. in order to speed up processing of Ipv6 packets by reducing the amount of information to be processed.

Regarding **claim 5**, Soirinsuo et al. and Conta et al. teach all of the limitations as described above with Soirinsuo et al. further teaching:

The method as defined in claim 1 wherein if no cache entry is found, performing a serial read of the extension headers and caching information on the these extension headers for processing subsequent packets in the same flow (column 8, lines 19-22, where a packet undergoes normal processing if a cache entry is not found, and column 7, line 64-column 8, line 14, where cache entries are formed based on extension headers).

Soirinsuo et al. does not teach using the lengths of the extension headers to assist in classifying packets. The general concept of using lengths of extension headers to process packets is well known in the art as illustrated by Conta et al. Conta et al. uses header length information to assist in processing the packet (page 26, section A.3). It would have been obvious to one of ordinary skill in the art at the time of invention to

modify Soirinsuo et al. with using header lengths as part of the classification as taught by Conta et al. in order to speed up processing of Ipv6 packets by reducing the amount of information to be processed.

Regarding **claim 6**, Soirinsuo et al. and Conta et al. teach all of the limitations as described above, with Soirinsuo et al. further teaching:
The method as defined in claim 1 including the step of detecting packets with hop-by-hop and routing extension headers and determining whether options processing of those packets is required (column 7, line 61-column 8, line 6, where the router determines what options to process including updating a hop-by-hop option or updating the routing header).

Regarding **claim 7**, Soirinsuo et al. and Conta et al. teach all of the limitations as described above, with Soirinsuo et al. not teaching:
The method as defined in claim 6 wherein, responsive to a packet not containing extension headers, read the upper-layer header without performing a cache lookup. The general concept of not performing a cache lookup on packets without extension headers is well known in the art as illustrated by Contra et al. Contra et al. teaches that classification of packets without extension headers can be performed without the need of a cache (page 14, section 5.7.2, where Ipv4 packets are seen as not having extension headers, and they are classified without a cache). It would have been obvious to one of ordinary skill in the art at the time of invention to modify

Soirinsuo et al. with not performing the cache lookup on packets without extension headers as taught by Conta et al. in order to increase reliability of the packet classification as the packet is processed correctly every time.

Regarding **claim 8**, Soirinsuo et al. and Conta et al. teach all of the limitations as described above, with Soirinsuo et al. further teaching:
The method as defined in claim 1 wherein if the extension headers are constantly changing resulting in incorrect cached data, the extension headers are traversed serially (column 7, line 64-column 8, line 22, where the extension headers are processed to create a cache entry, and then the cache entry is used to process subsequent packets that are similar to the cache entry, thus, if the cache entry has become incorrect, the extension headers must be different, and a new cache will be formed, this being seen as reading extension headers serially if the cache data is incorrect for the current flow).

Regarding **claims 9 and 10**, Soirinsuo et al. and Conta et al. teach all of the limitations as described above, with Soirinsuo et al. teaching:
The method as defined in claim 8 wherein the constantly changing extension headers are detected using a manual configuration and
The method as defined in claim 8 wherein the constantly changing extension headers are determined dynamically based on observations of the packet flow (column 7, lines 61-64, where routers can set up new flow handling states, seen as the ability to detect

changing extension headers, since the flow states are based on extension headers as seen in column 7, line 64-column 8, line 14, by observing a packet flow or by receiving a control protocol option to set up the flow information, seen as manually or dynamically setting up a flow handler, or detecting changing extension headers).

Regarding **claim 16**, Soirinsuo et al. and Conta et al. teach all of the limitations as described above, with Soirinsuo et al. further teaching:
The method as defined in claim 1 wherein results of classification and lookups performed during regular packet processing are cached (column 7, line 64-column 8, line 14, where the classification output is cached for subsequent packets to use).

Regarding **claim 19**, Soirinsuo et al. and Conta et al. teach all of the limitations as described above, with Soirinsuo et al. teaching that classifications of packets can be processed and cached and retrieved with cache lookups (column 8, lines 1-14). Soirinsuo et al. does not teach reading information from the header for reading the first extension header. The general concept of reading a header in the process of classification for reading a first extension header (which classification can be cached as Soirinsuo et al. teaches) is well known in the art as illustrated by Conta et al. Conta et al. teaches that the header length header is read to determine the length of the rest of the extension header (page 26, section A.3, where the header length header field is seen as an extension header read first, and this information assists the classifier to skip over the header's length simultaneously, or in parallel). It would have been obvious to

one of ordinary skill in the art at the time of invention to modify Soirinsuo et al. with reading extension header lengths as the first extension header in a classification system as taught by Conta et al. in order to speed up processing of Ipv6 packets by reducing the amount of information to be processed.

Regarding **claim 20**, Soirinsuo et al. and Conta et al. teach all of the limitations as described above, with Soirinsuo et al. further teaching:

The system as defined in claim 19 wherein if no cache entry is found, the extension headers are traversed serially (column 7, line 64-column 8, line 22, where the extension headers are processed to create a cache entry, and then the cache entry is used to process subsequent packets that are similar to the cache entry, thus, if the cache entry has become incorrect or is not found, the extension headers must be different, and a new cache will be formed, this being seen as reading extension headers serially if the cache data is incorrect for the current flow).

Regarding **claim 21**, Soirinsuo et al. and Conta et al. teach all of the limitations as described above, with Soirinsuo et al. further teaching:

The system as defined in claim 17 including means to specify, in certain situations, that cached data should not be used to attempt to accelerate packet processing (column 8, lines 15-24, where cached data expires and is not used for accelerated processing of packets).

4. **Claims 11-15 and 22-23** are rejected under 35 U.S.C. 103(a) as being unpatentable over Soirinsuo et al. (US 6,084,885) and "A proposal for the Ipv6 Flow Label" (Conta et al.) as applied to claims 1 and 17 above, and further in view of Zenchelsky et al. (US 6,173,364).

Regarding **claim 11**, Soirinsuo et al. and Conta et al. teach all of the limitations as described above, except for:

The method as defined in claim 1 wherein if the cached data does not match the packet the cache entry is not updated.

The general concept of not updating a cache entry if the cached data does not match the packet is well known in the art as illustrated by Zenchelsky et al. Zenchelsky et al. teaches a packet filtering system that uses a cache and each session entry is updated only if the session entry is used, and if the session entry is not used, the updating process of the session entry is not performed (column 5, line 55-column 6, line 4, where updating is seen as version checking and ejecting). It would have been obvious to one of ordinary skill in the art at the time of invention to modify Soirinsuo et al. and Conta et al. with using updating of cache entries only if the cache matches the packet, or not updating the entry if the cache data does not match the packet, as taught by Zenchelsky et al. in order to lessen the amount of times a cache is updated as noted in Zenchelsky et al.'s disclosure in column 4, lines 30-35.

Regarding **claims 12 and 13**, Soirinsuo et al., Conta et al., and Zenchelsky et al. teach all of the limitations as described above, with Soirinsuo et al. and Conta et al. not teaching:

The method as defined in claim 11 wherein detection that the cache does not match the packet is enabled manually or the method as defined in claim 11 wherein detection that the cache does not match the packet is enabled based on flow observations. The general concept of enabling detection that a cache does not match the packet manually or based on flow observations is well known in the art as illustrated by Zenchelsky et al. Zenchelsky et al. teaches a cache system for a packet filter that uses flow observations to enable the detection that the cache does not match the packet (column 5, lines 53-67, where the session entries are searched for a match, and this uses actual packet flow to enable the detection that a cache entry does not match the packet). Zenchelsky et al. also teaches manual enablement of the detection that a cache does not match the packet (column 5, lines 53-67, where even if the cache matches the packet, a change in a local rule base can cause an ejection of the cache entry, seen as not matching, and the local rule base update is seen as a manual enablement of detection that a cache does not match the packet). It would have been obvious to one of ordinary skill in the art at the time of invention to modify Soirinsuo et al., Conta et al., and Zenchelsky et al., with the further teaching of Zenchelsky et al. in order to lessen the amount of times a cache is updated as noted in Zenchelsky et al.'s disclosure in column 4, lines 30-35.

Regarding **claims 14-15 and 22-23**, Soirinsuo et al. and Conta et al. teach all of the limitations as described above except for:

The method as defined in claim 1 or the system of claim 17 wherein information from the upper-layer header is cached and the method as defined in claim 14 or the system of claim 22 wherein the cached information includes protocol and source and destination port identification. The general concept of caching upper layer header information including protocol and source and destination port identification is well known in the art as illustrated by Zenchelsky et al. Zenchelsky et al. teaches a system using a packet cache that caches a 5 tuple (column 4, lines 45-53, where a 5 tuple consists of protocol and source and destination ports as seen in column 1, lines 26-30). It would have been obvious to one of ordinary skill in the art at the time of invention to modify Soirinsuo et al. and Conta et al. with caching upper layer header information as taught by Zenchelsky et al. in order to reduce processing of similar packets as noted in Zenchelsky et al.'s disclosure in column 4, lines 9-11.

Conclusion

5. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure.

"Design of Multi-field Ipv6 Packet Classifiers Using Ternary CAMs" (Huang et al.) describes a hardware approach to Ipv6 packet classification.

Hughes et al. (US 5,842,040) describes a packet cache system for streaming protocol data units.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Adam S. Weintrop whose telephone number is 571-270-1604. The examiner can normally be reached on Monday through Friday 7:30am-5:00pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Jason Cardone can be reached on 571-272-3933. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

AW 6/12/07

JASON CARDONE
SUPERVISORY PATENT EXAMINER